Deep Learning-Based Approximation of Optimal Traveling Salesman Tour Length

Karlo Bala[a], Dejan Brcanov[b], Nebojša Gvozdenović[b], Andrea Rožnjik,[c]

[a]The Institute for Artificial Intelligence Research and Development of Serbia, Fruskogorska 1, 21000 Novi Sad, Serbia, karlo.bala@ivi.ac.rs
[b]Faculty of Economics, University of Novi Sad, Segedinski put 9-11, 24000 Subotica, Serbia, dejan.brcanov@ef.uns.ac.rs, nebojsa.gvozdenovic@ef.uns.ac.rs
[c]Faculty of Civil Engineering Subotica, University of Novi Sad, Kozaračka 2/A, 24000 Subotica, Serbia, andrea@gf.uns.ac.rs

This study introduces a deep learning approach to solve the Traveling Salesman Problem (TSP) in non-Euclidean spaces. Our main objective is to train a deep neural network to efficiently estimate the route length by understanding the problem's topology. We use a dataset of 14.4 million TSP instances obtained from real-world problems, previously solved using brute-force methods, covering problem sizes from 4 to 12 cities. To address this problem, we employed a five-layer multilayer perceptron with the ReLU activation function, trained for 4096 epochs. Our input data comprises a matrix representing the distances between cities. Depending on the problem size, our model achieved an average deviation from optimal solutions ranging from 2% to 2.9%. Additionally, the R^2 metric between predicted and target values is 0.992.

While this initial experiment shows promising results, it's just the first step toward estimating solutions for more complex vehicle routing problems.